

transactions. For each journal entry there is a corresponding TargetPage object that contains the Debits Target and Credits Target for that journal entry.

As the student manipulates the application interface, each action is reported to the ICAT. In order to tell the ICAT what actions were taken, the application calls to a database and asks for a specific interface control's ID. When the application has the ID of the target control and the SourceItem control, the application notifies the ICAT about the Target to SourceItem mapping. In other words, every time a student manipulates a source item and associates it with a target (e.g., dragging an account name to a debit line in the journal), the user action is recorded as a mapping of the source item to the target. This mapping is called a UserSourceItemTarget. Figure 21 illustrates the mapping of a source item to a target item in accordance with a preferred embodiment. When the student is ready, he submits his work to one of the simulated team members by clicking on the team member's icon. When the ICAT receives the student's work, it calculates how much of the work is correct by concept. Concepts in our journalization activity will include Debits, Credits, Asset Accounts, etc. For each of these concepts, the ICAT will review all student actions and determine how many of the student actions were correct. In order for the ICAT to understand which targets on the interface are associated with each concept, the targets are bundled into target groups and prioritized in a hierarchy. Once all possible coach topics are activated, a feedback selection analyzes the active pieces of remediation within the concept hierarchy and selects the most appropriate for delivery. The selected pieces of feedback are then assembled into a cohesive paragraph of feedback and delivered to the student. Figure 23 illustrates a feedback selection in accordance with a preferred embodiment. After the ICAT has activated CoachTopics via Rule firings, the Feedback Selection Algorithm is used to determine the most appropriate set of CoachItems (specific pieces of feedback text associated with a CoachTopic) to deliver. The Algorithm accomplishes this by analyzing the concept hierarchy (TargetGroup tree), the active CoachTopics, and the usage history of the CoachItems. Figure 24 is a flowchart of the feedback logic in accordance with a preferred embodiment. There are five main areas to the feedback logic which execute sequentially as listed below. First, the algorithm looks through the target groups and looks for the top-most target group with an active coach topic in it. Second, the algorithm then looks to see if that top-most coach item is praise feedback. If it is praise feedback, then the student has correctly completed the business deliverable and the ICAT will stop and return that coach item. Third, if the feedback is not Praise, then the ICAT will look to see if it is redirect, polish, mastermind or incomplete-stop. If it is any of these, then the algorithm will stop and return that feedback to the user. Fourth, if the feedback is focus, then the algorithm looks to the children target groups and groups any active feedback in these target groups with the focus group header. Fifth, once the feedback has been gathered, then the substitution language is run which replaces substitution variables with the proper names. Once the ICAT has chosen the pieces of feedback to return, the feedback pieces are assembled into a paragraph. With the paragraph assembled, the ICAT goes through and replaces all variables. There are specific variables for SourceItems and Targets. Variables give feedback specificity. The feedback can point out which wrong SourceItems were placed on which Targets. It also provides hints by providing one or two SourceItems which are mapped to the Target.

The Steps Involved in Creating Feedback in Accordance With A Preferred Embodiment

The goal of feedback is to help a student complete a business deliverable. The tutor needs to identify which concepts the student understands and which he does not. The tutor needs to tell the student about his problems and help him understand the concepts. There are seven major steps involved in developing feedback for an application. First, creating a strategy — The designer defines what the student should know. Second, limit errors through interface — The designer determines if the interface will identify some low level mistakes. Third, creating a target group hierarchy — The designer represents that knowledge in the tutor. Fourth, sequencing the target group hierarchy — The designer tells the tutor which concepts should be diagnosed first.

Fifth, writing feedback — The designer writes feedback which tells the student how he did and what to do next. Sixth, writing Levels of Feedback — The designer writes different levels of feedback in case the student makes the same mistake more than once. Seventh, writing rules — The designer defines patterns which fire the feedback.

A feedback strategy is a loose set of questions which guide the designer as he creates rules and feedback. The strategy describes what the student should learn, how he will try and create the business deliverable and how an expert completes the deliverable. The goal of the application should be for the student to transition from the novice model to the expert model. *What should the student know after using the application?* The first task a designer needs to complete is to define exactly what knowledge a student must learn by the end of the interaction. Should the student know specific pieces of knowledge, such as formulas? Or, should the student understand high level strategies and detailed business processes? This knowledge is the foundation of the feedback strategy. The tutor needs to identify if the student has used the knowledge correctly, or if there were mistakes. An example is the journal task. For this activity, students need to know the purpose of the journalizing activity, the specific accounts to debit/credit, and how much to debit/credit. A student's debit/credit is *not correct or incorrect* in isolation, but correct and incorrect in connection with the dollars debited/credited. Because there are two different types of knowledge—accounts to debit/credit and amounts to debit/credit—the feedback needs to identify and provide appropriate feedback for both types of mistakes.

How will a novice try and complete the task? Designers should start by defining how they believe a novice will try and complete the task. Which areas are hard and which are easy for the student. This novice view is the mental model a student will bring to the task and the feedback should help the student move to an expert view. Designers should pay special attention to characteristic mistakes they believe the student will make. Designers will want to create specific feedback for these mistakes. An example is mixing up expense accounts in the journal activity. Because students may mix up some of these accounts, the designer may need to write special feedback to help clear up any confusion.

How does an expert complete the task? This is the expert model of completing the task. The feedback should help students transition to this understanding of the domain. When creating feedback, a designer should incorporate key features of the expert model into the praise feedback he writes. When a student completes portion of the task, positive reinforcement should be provided which confirms to the student that he is doing the task correctly and can use the same process to complete the other tasks. These four questions are not an outline for creating feedback, but they define what the feedback and the whole application needs to accomplish. The designer should make sure that the feedback evaluates all of the knowledge a student should learn. In addition, the feedback should be able to remediate any characteristic mistakes the designer feels the student will make. Finally, the designer should group feedback so that it returns feedback as if it were an expert. With these components identified, a designer is ready to start creating target group hierarchies. Because there are positive and negative repercussions, designers need to select the when to remediate through the interface carefully. The criteria for making the decision is if the mistake is a low level data entry mistake or a high level intellectual mistake. If the mistake is a low level mistake, such as miss-typing data, it may be appropriate to remediate via the interface. If the designer decides to have the interface point out the mistakes, it should look as if the system generated the message. System generated messages are mechanical checks, requiring no complex reasoning. In contrast, complex reasoning, such as why a student chose a certain type of account to credit or debit should be remediated through the ICAT.

System messages - It is very important that the student know what type of remediation he is going to get from each source of information. Interface based remediation should look and feel like system messages. They should use a different interface from the ICAT remediation and should have a different feel. In the journalization task described throughout this paper,

there is a system message which states "Credits do not equal debits." This message is delivered through a different interface and the blunt short sentence is unlike all other remediation. The motivation for this is that low level data entry mistakes do not show misunderstanding but instead sloppy work. Sloppy-work mistakes do not require a great deal of reasoning about why they occurred instead, they simply need to be identified. High-level reasoning mistakes, however, do require a great deal of reasoning about why they occurred, and the ICAT provides tools, such as target groups, to help with complex reasoning. Target group hierarchies allow designers to group mistakes and concepts together and ensure that they are remediated at the most appropriate time (i.e., Hard concepts will be remediated before easy concepts). Timing and other types of human-like remediation require the ICAT; other low-level mistakes which do not require much reasoning include: **Incomplete-** If the task requires a number of inputs, the interface can check that they have all been entered before allowing the student to proceed. By catching empty fields early in the process, the student may be saved the frustration of having to look through each entry to try and find the empty one. **Empty-** A simple check for the system is to look and see if anything has been selected or entered. If nothing has been selected, it may be appropriate for the system to generate a message stating "You must complete X before proceeding". **Numbers not matching-** Another quick check is matching numbers. As in the journalization activity, is often useful to put a quick interface check in place to make sure numbers which must match do. Small data entry mistakes are often better remediated at the interface level than at the tutor or coach level (when they are not critical to the learning objectives of the course). There are two main issues which must be remembered when using the interface to remediate errors. First, make sure the interface is remediating low level data entry errors. Second, make sure the feedback looks and feels different from the ICAT feedback. The interface feedback should look and feel like it is generated from the system while the ICAT feedback must look as if it were generated from an intelligent coach who is watching over the student as he works.

Creating the Target Group Hierarchy- Target groups are sets of targets which are evaluated as one. Returning to the severity principle of the feedback theory, it is clear that the tutor needs to identify how much of the activity the student does not understand. Is it a global problem and the student does not understand anything about the activity? Or, is it a local problem and the student simply is confused over one concept? Using the feedback algorithm described earlier, the tutor will return the highest target group in which there is feedback. This algorithm requires that the designer start with large target groups and make sub-groups which are children of the larger groups. The ICAT allows students to group targets in more than one category. Therefore a debit target for transaction thirteen can be in a target group for transaction thirteen entries as well as a target group about debits and a target group which includes all source documents. Target should be grouped with four key ideas in mind. Target groups are grouped according to: Concepts taught; Interface constraints; Avoidance of information overload and Positive reinforcement.

The most important issue when creating target groups is to create them along the concepts students need to know to achieve the goal. Grouping targets into groups which are analogous to the concepts a student needs to know, allows the tutor to review the concepts and see which concepts confuse the student. As a first step, a designer should identify in an unstructured manner all of the concepts in the domain. This first pass will be a large list which includes concepts at a variety of granularities, from small specific concepts to broad general concepts. These concepts are most likely directly related to the learning objectives of the course. With all of the concepts defined, designers need to identify all of the targets which are in each target group. Some targets will be in more than one target group. When a target is in more than one target group, it means that there is some type of relationship such as a child relationship or a part to whole relationship. The point is not to create a structured list of concepts but a comprehensive list. Structuring them into a hierarchy will be the second step of the process.